arXiv:1704.08364v1 [cs.DC] 26 Apr 2017

Low-complexity Distributed Tomographic Backprojection for large datasets

Gilberto Martinez Jr.*[1], Janito V. Ferreira Filho[1], and Eduardo X. Miqueles†[1]

[1] *Brazilian Synchrotron Light Laboratory (LNLS)/CNPEM, Campinas, SP, Brazil*

Abstract

In this manuscript we present a fast GPU implementation for tomographic reconstruction of large datasets using data obtained at the Brazilian synchrotron light source. The algorithm is distributed in a cluster with 4 GPU's through a fast pipeline implemented in C programming language. Our algorithm is theoretically based on a recently discovered low complexity formula, computing the total volume within $O(N^3 \log N)$ floating point operations; much less than traditional algorithms that operates within $O(N^4)$ *flops* over an input data of size $O(N^3)$. The results obtained with real data indicate that a reconstruction can be achieved within 1 second provided the data is transferred completely to the memory.

1 Introduction

In this manuscript, we present a fast implementation of the well-known *filtered-backprojection algorithm* (FBP) [1, 2, 3], which has the ability to reproduce reliable image reconstructions in a reasonable amount of time, before taking further decisions. The FBP is easy to implement and can be used to take fast decisions about the quality of the measurement, i.e., sample environment, beam-line conditions, among others. Figure 1 shows the fluxogram for an ideal tomographic experiment. A synchrotron facility able to measure a three-dimensional dataset $\boldsymbol{Y}$ within few seconds, needs a fast reconstruction algorithm able to provide a fast "preview" of the tomography within the same amount of time (Fig.1.(a.1)). If the experimental conditions are not satisfactory, the quality of the reconstruction will decrease, and the researcher can decide either to make another scan, or to process later the data using advanced reconstruction algorithms or even high quality segmentation methods (Fig.1.(d)). The difficulty here is that the FBP algorithm consist basically in two mathematical operators, which are filtering and backprojection. Filtering is FFT based since is a low-pass convolution operation [3]. Backprojection, on the other hand, is defined as an average through all the x-rays passing at a given pixel; therefore presenting high computational complexity of $O(N^3)$ for an image of N^2 pixels. The brute-force approach to compute the backprojection operator can be made extremely slow, even using a GPU implementation. Sophisticated ray-tracing strategies can also be used to make the running time faster and others analytical strategies reduce the backprojection complexity to $O(N^2 \log N)$, see [4, 5]. Our approach to compute the backprojection is called BST - as an acronym to *backprojection-slice theorem* [6] - having the same low complexity of $O(N^2 \log N)$ although easier to implement than his competitors, producing less numerical artifacts and following a more traditional "gridding strategy" similar as [7]. There are several others reconstruction softwares reproducing quasi-real time reconstructions, [8, 9, 10, 11].

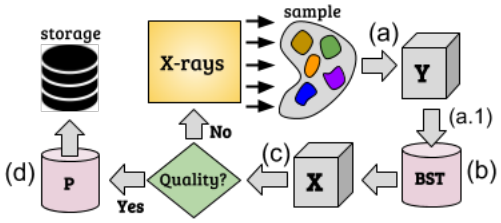

Figure 1: Fluxogram for a fast tomographic scan.

The computational gain of BST over the brute-force approach for computing the backprojection operator is presented using a fast pipeline for the data access. Our results indicate that reconstructions

*FAPESP 2016/16238-4
†CNPQ 442000/2014-6

through BST can be performed within 1 second using the IBM/MINSKY (4 NVIDIA P100) for datasets of size $2048 \times 2048 \times 2048$. A comparison is made with a SGI/TESLA using 4 NVIDIA K80. To make a reliable comparison, all codes were implemented without taking advantages of the NVIDIA-NVLINK. Also, we have tested our algorithm at small dimensions using only one GPU. Devices like JETSON TX1, indicates that a fast reconstruction is possible for images of size $512 \times 512 \times 512$, a conventional GPU like GT740M usually coupled with domestic notebooks can handle datasets of size $1024 \times 1024 \times 1024$ and finally a TITAN-X coupled with a standard PC handle volumes of size $2048 \times 2048 \times 2048$. In this sense, the reconstruction package supply three different aspects of a tomographic experiment: i) live reconstruction at the beamline assisting fast decisions by the researcher, ii) advanced reconstructions after the measurement using the beamline GPU power iii) reconstructions of the dataset without the beamline GPU power, at conventional desktops/notebooks.

Description of the problem: We consider the transmission tomographic problem using x-rays generated at synchrotron facility based on parallel x-rays. The Radon transform has been used extensively as the mathematical object modelling the inverse problem. A typical tomographic measure provides a three-dimensional dataset $\mathbf{Y}$ as presented in Figure 2.a. Dataset $\mathbf{Y}$ is such that each plane $s \times t$ determines a radiography of the sample, for a constant projection angle θ_j varying on a discrete mesh with V points from 0 to π. The plane $s \times t$ is a discretization of a CCD camera with $N \times N$ pixels. The pair (N, V) is a characteristic of the imaging beamline, typically $N \sim 2048$ and $V > 1001$, which means that $\mathbf{Y}$ is a large dataset with approximately 9GiB. The tomographic problem for parallel rays is posed

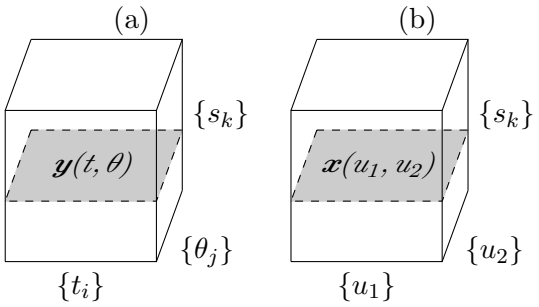

Figure 2: Definition of three-dimensional datasets: (a) Measured $\mathbf{Y}$ (b) Reconstructed $\mathbf{X}$.

in the following manner: find a reconstructed three-dimensional dataset $\mathbf{X}$ from $\mathbf{Y}$ in such a way that a given slice (s_k constant) $\boldsymbol{x}(u) = \boldsymbol{x}(u_1, u_2)$ of the cube $\mathbf{X}$ is related to the same slice $\boldsymbol{y}(t, \theta)$ from volume $\mathbf{Y}$ through the linear operation

$$\boldsymbol{y}(t,\theta) = \mathscr{R}\boldsymbol{x}(t,\theta) = \int_{\mathbb{R}^2} \boldsymbol{x}(u)\boldsymbol{\delta}(t - u \cdot \xi_\theta)\mathrm{d}u \qquad (1)$$

with $\xi = (\cos\theta, \sin\theta)$. Equation (1) is the Radon transform [3] from the two-dimensional function $\boldsymbol{x}$. Inverting the operator $\mathscr{R}$ is the mathematical core of most tomographic problems. There are several numerical algorithms for this task. One of the most celebrated algorithms is known as *filtered-backprojection*, given by the following formula $\boldsymbol{x} = \mathscr{B}[\mathcal{F}\boldsymbol{y}]$ where $\mathscr{B}$ is the *backprojection operator* - the adjoint of $\mathscr{R}$ - defined as the following integral operator

$$\boldsymbol{b}(u) = \mathscr{B}[\boldsymbol{y}](u) = \int_0^\pi \boldsymbol{y}(u \cdot \xi_\theta, \theta)\mathrm{d}\theta \qquad (2)$$

The operator $\mathcal{F}$ is a convolution, acting only on the first variable t, i.e., $\boldsymbol{h}(t,\theta) = \mathcal{F}\boldsymbol{y}(t,\theta) = \boldsymbol{y}(t,\theta) \star \ell(t)$ with $\hat{\ell}(\sigma) = |\sigma|$. Figure 3 illustrate the filtered-backprojection action, together with $\mathscr{B}$, $\mathscr{R}$ and $\mathcal{F}$ on a point source function $\boldsymbol{x}(u) = \boldsymbol{\delta}(u-a)$ (for a random point $a \in \mathbb{R}^2$). The function $\boldsymbol{y} = \mathscr{R}\boldsymbol{x}$ is often referred to as a sinogram because the Radon transform of an off-center point source is a sinusoid. The backprojection operation simply propagates the measured sinogram back into the image space along the projection paths. Under physical assumptions that are beyond

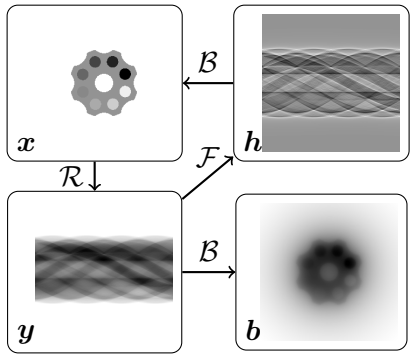

Figure 3: Action of operators Radon $\mathscr{R}$, Backprojection $\mathscr{B}$ and Filtering $\mathcal{F}$.

the scope of this manuscript, the photon propagation through a sample obeys the Lambert-Beer law $\boldsymbol{I}(\eta) = \boldsymbol{I}_0(\eta)e^{-\boldsymbol{y}(\eta)}$, $\eta = (t, \theta)$ with $\boldsymbol{I}, \boldsymbol{I}_0$ standing for the transmitted and incident photon counting on the pixel camera, parameterized by the point (t, θ)

and over the slice s_k. The filtered backprojection algorithm applied over the sinogram $\boldsymbol{y}$ (FBP) offers a good reconstruction under strictly severe conditions on the measured data $\{\boldsymbol{I}, \boldsymbol{I}_0\}$, almost never satisfied at real measurements. There are three main processing steps on the data, before reconstruction, they are: **(a)** *Normalization*: the sinogram $\boldsymbol{y}$ is obtained theoretically using the logarithmic function. In practice, a dark/flat field correction is used $\boldsymbol{y} = -\log\left[(\boldsymbol{I} - \boldsymbol{D})/(\boldsymbol{I}_0 - \boldsymbol{D})\right]$, where $\boldsymbol{D}$ is the dark-field measurement (i.e., without the sample). **(b)** *Centering*: In practice, the tomographic device suffers from several mechanical imprecisions, most of them, intrinsic to the problem. In this case, we obtain an experimental sinogram which is a not an exact realization of a Radon transform. This is case of a sample rotating above a precision stage, which is not perfectly aligned to the camera vertical axis. Here, the center of rotation of the sample is unknown, and the originated sinogram is slightly shifted from the correct center of rotation. As a consequence, the measured sinogram $\boldsymbol{y}$ is a noisy and shifted version of the theoretical one $\boldsymbol{y}(t - \beta, \theta)$ where β is the unknown shift.
(c) *Ring Filtering*: Dead camera pixels or imperfections on the scintillator remain constant as the angle θ varies within the interval $[0, \pi]$, giving origin to constant artifacts on the gathered frames $\boldsymbol{I}$. Therefore, strong stripes arise in the sinogram $\boldsymbol{y}$, producing concentric rings on the final reconstructed image. There are also several ring artifacts correction algorithms, producing an approximation of $\boldsymbol{y}$ as discussed in [12].

2 Fast Backprojection

The backprojection corresponds to an average 'smear' of all projections passing through a single pixel $\boldsymbol{x}$. Since an average is described as a convolution, it is natural to regard formula (2) as an average in the frequency domain. It was shown recently [6] that the operator $\mathscr{B}\boldsymbol{y}$ can be computed within $O(N^2 \log N)$ flops per slice, assuming that $\boldsymbol{y}$ is a sinogram having dimension of order $O(N^2)$. This is not the only fast approach to compute $\mathscr{B}$, many others can be found in [5, 4] while other algorithms as [7, 13] propose to reconstruct $\boldsymbol{x}$ using $\boldsymbol{y}$ as an input, through the well known *Slice Theorem* or *Fourier-Slice Theorem* (FST) relating the values of $\boldsymbol{y}$ and $\boldsymbol{y}$ in the reciprocal space. Our approach, the so-called *Backprojection-slice Theorem* (BST) is based on the following theorem:

Theorem 2.1. *Let $\boldsymbol{y}$ be a given sinogram and $\widehat{\cdot}$ denotes the Fourier transform operation. The backprojection $\boldsymbol{b} = \mathscr{B}\boldsymbol{y}$ can be computed through*

$$\widehat{\boldsymbol{b}}(\sigma\cos\theta, \sigma\sin\theta) = \hat{\boldsymbol{y}}(\sigma, \theta)/\sigma \tag{3}$$

with $\sigma > 0 \in \mathbb{R}$ and $\theta \in [0, 2\pi]$.

The BST approach is the dual of the FST, in the sense that $\boldsymbol{y} = \mathscr{R}\boldsymbol{y}$ and $\boldsymbol{b} = \mathscr{B}\boldsymbol{y}$ are related in the reciprocal space within a polar grid. If $\boldsymbol{y} = \mathscr{R}\boldsymbol{x}$ the computation of $\boldsymbol{b} = \mathscr{R}\boldsymbol{x}$ follows easily convolving the feature image $\boldsymbol{x}$ with the point-spread function $1/\|u\|_2$ [3], but $\boldsymbol{x}$ is unknown for pratical experiments. Also, FST provide a direct link between $\{\boldsymbol{x}, \boldsymbol{y}\}$ and $\{\boldsymbol{b}, \boldsymbol{y}\}$. If $\boldsymbol{y} \neq \mathscr{R}\boldsymbol{x}$ the link between $\{\boldsymbol{x}, \boldsymbol{y}\}$ is no longer valid through FST but we can still provide a connection between $\{\boldsymbol{b}, \boldsymbol{y}\}$. This is the main goal of the BST approach. The straight usage of the Fourier transform for the implementation of (3) produce big artifacts near the origin, caused by the fact that the values on sinogram on the line $t = 0$ are not equal to 0. This problem can be solved with usage of short-time Fourier transform near the origin, with window w, e.g. Kaiser-Bessel window function. The BST strategy applied over a sinogram image $\boldsymbol{y}$ is obtained after 7 processing stages $\{P_k\}$, i.e. $\boldsymbol{b} = P_6P_5P_4P_3P_2P_1P_0(\boldsymbol{y})$.

Each processing step P_k can be implemented in a parallel form with CUDA. Step P_0 and P_1 indicates an interpolation to polar coordinates with the multiplication of the Kaiser-Bessel window function, respectively. This is an easy process, computed with complexity $O(NV)$. Step P_2 is the zero padding of the polar sinogram - equivalent to an oversampling in the frequency domain - with the same complexity of P_0. Even though P_1 and P_2 can be merged into one single step, they were considered disjoint operations in our customized implementation. Step P_3 is a one-dimensional Fourier Transform of the data from step P_2. Step P_4 is the convolution of the polar sinogram with kernel $1/\sigma$. This part was divided in m parallel FFT's, each computed at an individual thread using advanced strategies with complexity $O(N \log N)$. Step P_5 is an interpolation from polar to cartesian coordinates in the frequency domain. The bigger the zero padding at step P_2, the better this part will behave, preventing aliasing artifacts. Step P_6 is a two-dimensional inverse Fourier transform of the data from step P_5. This is an operation with

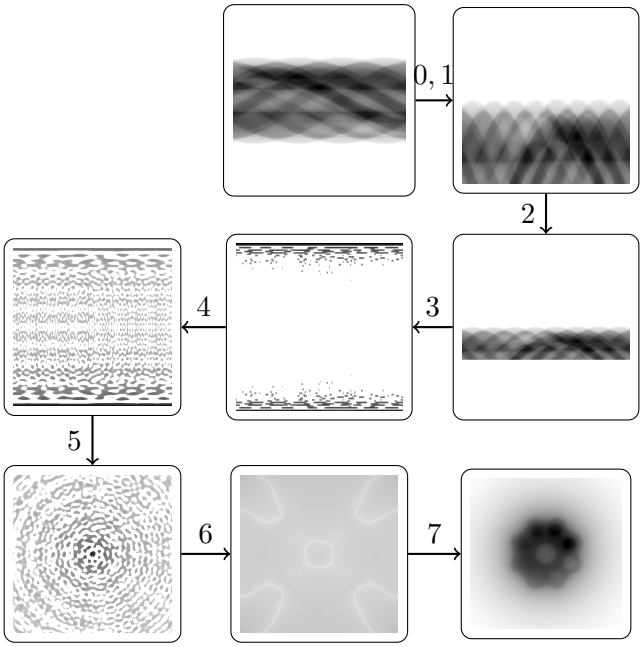

Figure 4: Theoretical fluxogram BST formula, from Eq.(3).

low computational complexity and obtained with order $O(N^2 \log N)$. Step P_7 is the FFT shifting from previous step. We could add a P_8 step, an optional two-dimensional interpolation of the resulting image to the correct feature domain $[-1, 1] \times [-1, 1]$. A theoretical fluxogram for BST, describing processing steps $\{P_0, \ldots, P_7\}$ is presented in Figure 4.

Computing the backprojection operator $\mathscr{B}$ directly from equation (2) produce a high-computational complexity algorithm which is easy to implement in a parallel structure, either in CPU or GPU. Such approach has a computational complexity of $O(N^3)$ per slice, which in turns implies a total cost of $O(N^4)$ for a set of N sinogram images N. We denote this as a *Slant-stack* approach (SS) since summation is performed over straight lines with slope according to the input angle θ [14]. A block of Q sinograms is usually processed at once due to the fact that the backprojection kernel can be processed simultaneously for a tomographic scan using parallel rays. The fluxogram presented in Figure 4 indicates that each step of BST has computational complexity bounded by $O(N^2 \log N)$.

3 Pipeline for Distribution

The full implementation is composed of multiple different stages: sinogram normalization (N), low-pass filtering (F), backprojection (B), ring filtering (R), centering of sinograms (R) and saving reconstructed data to storage (S). A typical sequence for these operations is presented in Figure 5. Although all stages operate on the same data, each has different characteristics regarding resource usage. Loading the input data and storing the output results are I/O bound operations, other stages are CPU bound and some are GPU bound. Even among the GPU bound stages, some are more compute intensive while others are more memory intensive.

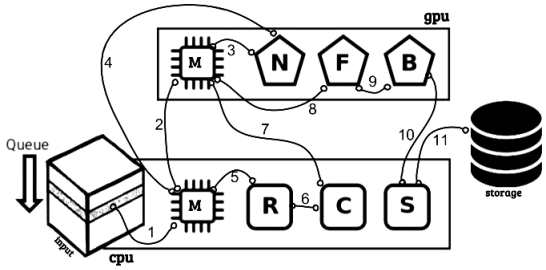

Figure 5: Pipeline E(CPU/GPU) using a generic backprojection kernel **B**.

Having such a heterogeneous set of process stages, it is desirable to optimize their execution in order to maximize resource utilization and consequently minimize execution time. Such optimization becomes non-trivial as the available resources increase and the scheduling possibilities are amplified. Although the stages must be executed in a given order for a given set of input data, they can be executed in parallel with different sets of input data, as long as each data set is processed by each stage in the expected order. To help with the scheduling of the available jobs (the execution of one of the stages on one of the input data sets), a helper framework was implemented. This framework consists of an abstract representation of a pipeline of interconnected stages. Data enters the pipeline and is processed by each stage sequentially until it reaches the end of the pipeline, producing the resulting output data for the given input data. Each stage is represented as a processing function that receives a generic data blob and returns a resulting generic data blob. The function therefore processes a data block into another data block. The contents of the data block is specific to each stage, and it is a requirement that the resulting output block is in the expected format for the input block of the stage that follows. With the implementation modeled as a pipeline, some resource optimizations can be applied

generically. The first is the execution of each stage in a separate CPU thread. This is not an optimization to primarily use different CPU cores to execute stages in parallel, although it does contribute to the usage of more resources. Instead, the main advantage is that it allows the operating system to schedule the different stages depending on their resource usage patterns. This means that I/O bound stages can start, and when an I/O operation halts the execution, a stage that is GPU bound can be executed, and when that stage is waiting on the results from the GPU another stage can be executed or the idle stages can be resumed if the data they're waiting for is ready.

	Jetson TX1		GT740M		Titan X		
DATASET	SS	BST	SS	BST	SS	BST	Q
2048^3					895.072083	190.321701	5
					799.828003	228.629852	10
1536^3					174.356842	56.577358	5
					253.741974	71.588837	10
1024^3			327.243317	62.341717	30.724861	17.413599	5
			259.171601	57.566849	24.224422	17.024681	10
512^3	26.321959	13.171736	23.42721	9.229262	3.858362	4.774002	5
	24.451975	12.754734	18.632298	8.237438	3.802566	3.744505	10
	21.064373	13.432063	16.567259	7.892296	4.142621	3.781465	15
256^3	3.516305	3.50024	1.948982	1.886878	1.022957	1.032447	5
	3.128555	3.290957	1.692543	1.437727	0.965807	0.969846	10
	3.041064	3.131222	1.567504	1.355728	0.965546	0.978986	15
2048^2x1024					1986.611572	274.715823	5
					1172.36645	517.887451	10
1536^2x768					72.134781	70.911224	5
					66.979645	28.303719	10
1024^2x512			165.145828	42.321434	14.960188	9.268951	5
			130.446228	35.056644	14.851658	9.82362	10
512^2x256	14.061594	8.30939	16.655327	5.380389	2.084591	2.029896	5
	13.922671	8.293776	9.346146	5.003119	1.998881	2.203163	10
	16.117891	8.202461	8.63123	4.846304	1.995749	1.907305	15
256^2x128	2.515128	2.740289	1.0246685	1.205075	0.643044	0.681087	5
	2.421637	2.650361	1.036608	1.056225	0.507826	0.525796	10
	2.237847	2.310273	0.957721	1.028077	0.469525	0.500501	15

Table 1: Execution time (in seconds) for the pipeline E(GPU), using only one GPU, and comparing back-projection kernels SS and BST.

With each stage executing in different threads, the jobs are executed concurrently. An optimization opportunity arises in managing the number of jobs in flight. In one extremity, the number of jobs in flight can be limited to one, which effectively makes the pipeline run sequentially with no concurrency. On the other extremity, the number can be unlimited. While this might appear to be the optimal solution, because generally stages don't have the same execution time, some stages end up being bottlenecks. When this occurs, many jobs remain queued before the slowest stage, and they can potentially be holding resources that they are not using while idle. An example of this was that between two GPU specific stages, the data transferred between them resided in the GPU memory. Because the second stage was slower than the first stage, the GPU ran out of memory due to the amount of queued jobs for the second stage. The solution to this problem was to allow limiting the number of jobs in queue before each stage. Each stage can have its maximum queue length individually tuned, to allow limiting only the jobs in flight that may strain resources while idle.

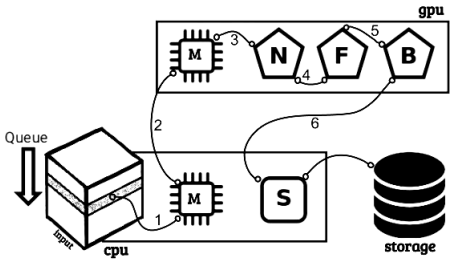

Figure 6: Pipeline for GPU code using a generic back-projection kernel **B**.

The final generic optimization is to allow running the same stage in a pool of threads. If the execution isn't fully using a specific resource, the stage that uses it most can be executed in more than one thread, which means they use that resource concurrently, potentially increasing its utilisation. An example is CPU core usage. If not all of the systems core are in use, have the same stage execute in more than one thread allows the operating system to schedule more threads on the CPU, attempting to maximize its usage. This also makes it easier when using multiple GPUs, since each thread can be assigned to work with a single GPU. Even having more than one thread per GPU can increase throughput, because even if the task is the same, the GPU driver can schedule data transfers between the main memory and the GPU memory while the GPU is actively running another job (this is allowed in CUDA as long as the program is configured to have a per-thread default stream).

Such abstract implementation of a work pipeline allows the algorithm to be broken up into stages - as shown in Figure 5 - and makes it easier to manage the execution of these stages in order to maximize resource usage. After adapting the implementation to the pipeline, some execution parameters can be tweaked to increase throughput. These include the number of threads for each stage, and the maximum queue size before each stage. This allows a simple way to experiment with different implementations by adapting the work done in each stage, and also quick experimentation with job scheduling to validate how different allocations result in different resource usage

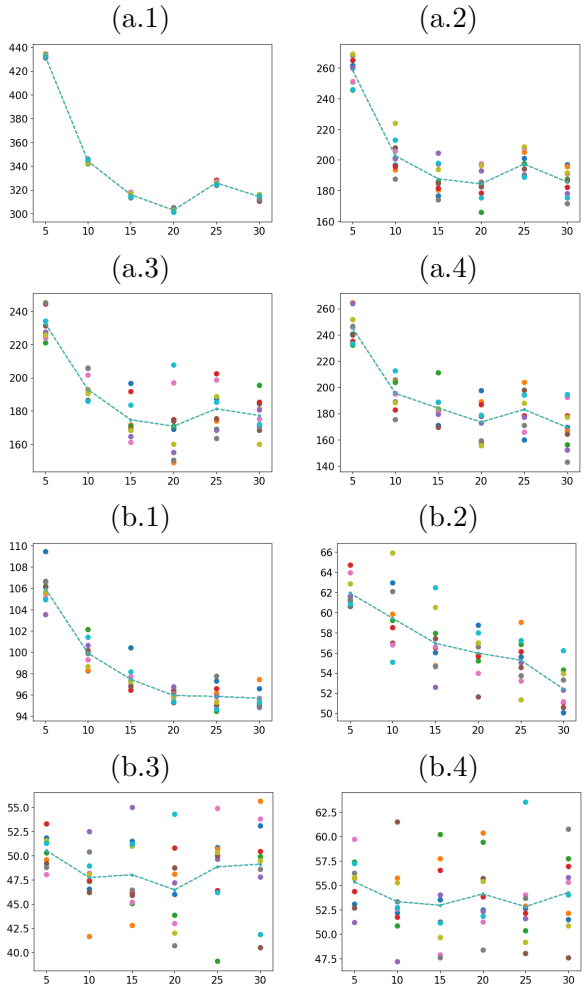

Figure 7: Execution time (in seconds) for pipeline E(GPU) as a function of blocksize Q, using 1,2,3,4 GPUs and the high-complexity kernel SS. (a) SGI/C2108-GP5 (b) Minsky/IBM.

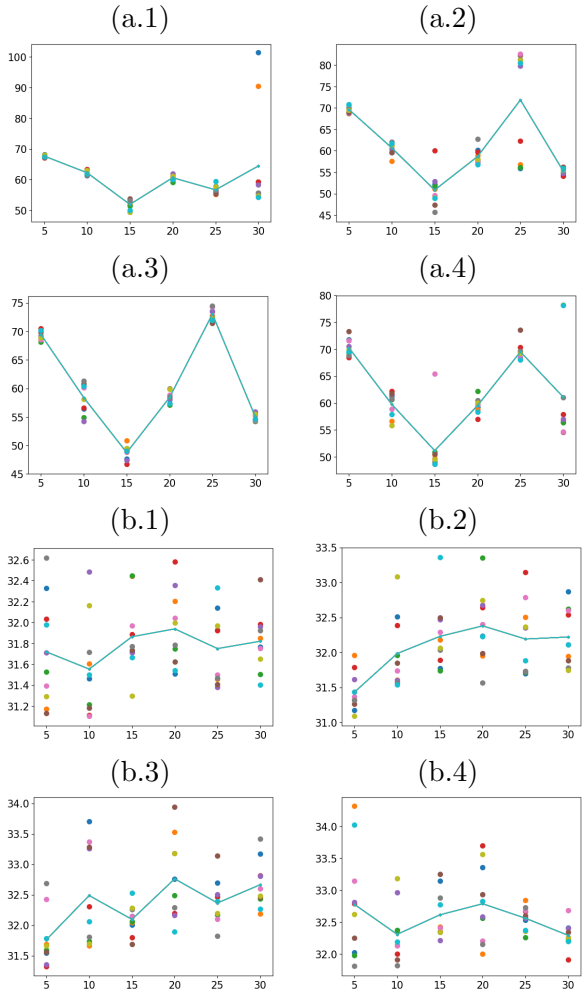

Figure 8: Execution time (in seconds) for pipeline E(GPU) as a function of blocksize Q, using 1,2,3,4 GPUs and the low-complexity kernel BST. (a) SGI/C2108-GP5 (b) Minsky/IBM

and performance.

JETSON TX1		GT740M			TITAN X			
256^3	512^3	256^3	512^3	1024^3	256^3	512^3	1024^3	2048^3
1.6	6.80	0.56	2.34	10.52	0.69	3.02	14.02	110.81

Table 2: Average time (in seconds) expended reading the input data $\boldsymbol{Y}$ using a non-optimized HDF call. All executions were done using blocksize of $Q = 10$.

4 Performance

Our pipeline, as presented in Section 3, is dependent on the time for reading the input data $\boldsymbol{Y}$ - see definition on Figure 2. The hierarchical data format HDF5 is becoming very popular for the storage of measured data at synchrotron facilities. In this sense, our algorithm is evaluated in two ways: **(a)** Extracting - without MPI capabilities - a block of Q images. Here, we have used the standard call of HDF to read the given number of images. The block of images is then processed following the pipeline of Figures 5 and 6, except that we are not concerned with output. Hence, the flux of reconstructed data to storage is not considered in our analysis. Executions using the pipeline of Figure 5 are refered as E(CPU/GPU) as they have an extra processing of parallel kernels **R** and **C** at CPU. Executions using the pipeline of Figure 6 are refereed as E(GPU). The pipeline was executed using $W = 8$ work items, each using $T = 2{\times}G$ threads, being G the number of availables GPUs. Needless to say that Q is bounded by the GPU memory and therefore by G and W, i.e., $A_{\mathbf{N}}(QWT) + A_{\mathbf{F}}(QWT) + A_{\mathbf{B}}(QWT) \leq \mathbf{M}$ with $A_{\mathbf{N}}$ the amount of memory allocation for kernel **N** and **M** the global memory of a single GPU. **(b)** Measuring the elapsed time of the pipeline just

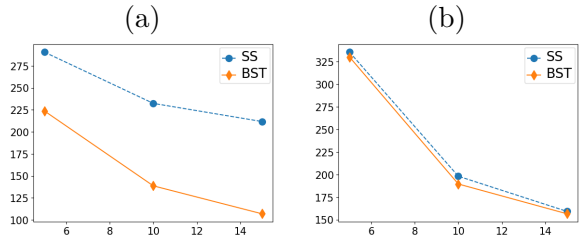

Figure 9: Execution time (in seconds) for pipeline E(CPU/GPU) as a function of blocksize Q, using 4 GPUs (a) SGI/C2108-GP5 (b) Minsky/IBM.

reading the input data $\mathbf{Y}$, and turning off all subsequent kernels. In average, we understand that the execution time of the reconstruction is bounded by this step of the pipeline. We denote this execution as R(CPU/GPU).

Table 1 presents the elapsed times for the pipeline E(GPU) using only one GPU and standard tomographic dimensions[1]. Using BST with a powerful GPU like TITAN X we can achieve a feasible reconstruction of a dataset of dimension 2048^3 within 3 minutes. At the same dimension and after 10 executions of the pipeline R(CPU/GPU) with a blocksize of $Q = 10$, we obtain an average of 107 seconds wasted on reading the input data - see Table 2. Hence, the reconstruction can be achieved within one minute. Using the same reasoning for the other devices, the reconstruction with JETSON runs within 6 seconds for a dimension of 512^2 and within 1 second for dimensions $512^2 \times 256$. The standard GPU-GT740 runs a reconstruction within 46 seconds for a dimension 1024^3.

	MINSKY/IBM						SGI/C2108-GP5					
Q	5	10	15	20	25	30	5	10	15	20	25	30
1GPU	30.3	30.10	30.46	30.38	29.92	29.85	38.70	41.81	38.45	42.47	39.41	41.77
2GPUs	30.37	29.97	30.29	30.16	30.29	30.21	39.13	40.89	40.96	38.95	39.21	41.87
3GPUs	30.70	30.37	30.51	30.37	30.60	30.30	35.92	37.73	40.54	39.66	41.94	42.45
4GPUs	30.58	30.75	30.55	30.67	30.20	30.61	35.99	37.52	42.43	39.68	42.88	42.15

Table 3: Average time (in seconds) expended reading the input data $\boldsymbol{Y}$ using a non-optimized HDF call, different blocksizes and two clusters.

A large number of executions (at two different clusters) of the pipeline E(GPU) using real data with dimensions $(N, V) = (2048, 2000)$ is presented in Figures 7 and 8. Plots (a.k) therein represents 10 elapsed times (in seconds) versus the blocksize number Q - with k GPUs for code distribution running at a SGI/C2108-GP5 server coupled with 4 NVIDIA K80. Plots (b.k) represents 10 elapsed times (in seconds) versus the blocksize number Q - with k GPUs for code distribution running at a Minsky/IBM server coupled with 4 NVIDIA P100. Times for BST (Fig.8) are much lesser than SS (Fig.7), as predicted by theory. The average time expended reading the input data for these servers using a non-optimized HDF call, is presented in Table 3 for different values of blocksize Q. The processing time expected for BST is near 1 second using 1,2,3 or 4 GPUS at Minsky cluster while 30 seconds for the SGI cluster. The average time for a complete reconstruction using the high-complexity kernel SS depends on the number of GPUs used for distribution. The elapsed times for the hybrid pipeline E(CPU/GPU) - see Figure 5 - running on the same data is presented in Figure 9. Now, it is clear that the time expended processing the data on the CPU affect performance dramatically. In fact, the CPU kernel **R** process a single slice using a Conjugate gradient method for stripe suppression [12]. Processing a block of Q images, although easy to implement, does not present the same quality of sinogram restoration because correction is slice dependent. Hence, even with a $O(N)$ computational complexity, the ring suppression plus data transfer from CPU/GPU should be avoided in this pipeline. There is certainly a huge space for improvements on the ring suppression algorithm, handling a block of sinograms. The CPU kernel **C** - for centering sinograms - is optional and easy to implement for a block of images not affecting the final performance.

5 Conclusions

Fast reconstructions of large tomographic datasets is feasible using the BST formula (3). The superscalar pipeline - running only GPU kernels with computational complexity bounded above by $O(N^2 \log N)$ - is able to complete a 3D reconstruction within 1 second of running time provided the input dataset is transferred completely to the memory. New imaging detectors capable to store the data in a local buffer could benefit from this reconstruction approach. We emphasize that other low-complexity algorithms [4, 8, 9] are also capable to provide fast re-

[1]We have adopted a three-dimensional dataset described by an ellipsoid. The feature image is defined as $\boldsymbol{x}(u) = \rho$ if $(u_1/A)^2 + (u_2/B)^2 \leq 1 - (s/C)^2$ and zero otherwise. Here, A and B are parameters defining the shape of the semi-major axis u_1, u_2 of the ellipse and C defining the length of slice axis $|s| \leq C$. The Radon transform of $\boldsymbol{x}$ can be obtained analytically [1]. A discretization of the sinogram using $\{N, V\}$ points for t and θ axis respectively, gives us a three-dimensional dataset $\boldsymbol{Y}$.

sults. Nonetheless, a fast backprojector (BST) and a fast projector (FST) implemented in the same pipeline presented here, give us the chance to implement advanced reconstruction strategies [15] (mostly iterative) with the same low computational complexity. Further improvements on the code[2] are under progress which include: usage of MPI strategies to read the data, ring-filtering kernel to process a block of sinograms, thread optimization among others.

Acknowledgments

We would like to thank the IBM team for providing access to the Poughkeepsie Benchmarking Technical Computing Cloud: Douglas M. Dreyer, Victoria Nwobodo, James Kuchler, Khajistha Fattu, Antonio C.Navarro and Leonardo A.G. Garcia. Thanks also to Harry Westfahl Jr. for many valuable suggestions. The TITAN X used for this research was donated by the NVIDIA Corporation.

References

[1] Avinash C Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*. IEEE press, 1988.

[2] Frank Natterer. *The mathematics of computerized tomography*, volume 32. Siam, 1986.

[3] Stanley R Deans. *The Radon transform and some of its applications*. Courier Corporation, 2007.

[4] Fredrik Andersson. Fast inversion of the radon transform using log-polar coordinates and partial back-projections. *SIAM Journal on Applied Mathematics*, 65(3):818–837, 2005.

[5] Ashvin George and Yoram Bresler. Fast tomographic reconstruction via rotation-based hierarchical backprojection. *SIAM Journal on Applied Mathematics*, 68(2):574–597, 2007.

[6] Eduardo Miqueles and Elias S Helou. Fast backprojection operator for synchrotron tomographic data. In *European Conference on Mathematics for Industry*. Springer, 2014.

[7] F Marone and M Stampanoni. Regridding reconstruction algorithm for real-time tomographic imaging. *Journal of synchrotron radiation*, 19(6):1029–1037, 2012.

[8] Doga Gürsoy, Francesco De Carlo, Xianghui Xiao, and Chris Jacobsen. Tomopy: a framework for the analysis of synchrotron tomographic data. *Journal of synchrotron radiation*, 21(5):1188–1193, 2014.

[9] Federica Marone, Alain Studer, Heiner Billich, Leonardo Sala, and Marco Stampanoni. Towards on-the-fly data post-processing for real-time tomographic imaging at tomcat. *Advanced Structural and Chemical Imaging*, 3(1):1, 2017.

[10] Roman Shkarin, Evelina Ametova, Suren Chilingaryan, Timo Dritschler, Andreas Kopmann, Alessandro Mirone, Andrei Shkarin, Matthias Vogelgesang, and Sergey Tsapko. Gpu-optimized direct fourier method for on-line tomography. *Fundamenta Informaticae*, 141(2-3):245–258, 2015.

[11] A Mirone, R Wilcke, A Hammersley, and C Ferrero. Pyhst–high speed tomographic reconstruction, 2010.

[12] Eduardo X Miqueles, Jean Rinkel, Frank O'Dowd, and JSV Bermúdez. Generalized titarenko's algorithm for ring artefacts reduction. *Journal of synchrotron radiation*, 21(6):1333–1346, 2014.

[13] Daniel Potts and Gabriele Steidl. New fourier reconstruction algorithms for computerized tomography. In *International Symposium on Optical Science and Technology*, pages 13–23. International Society for Optics and Photonics, 2000.

[14] A Averbuch, RR Coifman, DL Donoho, M Israeli, and J Walden. *Fast Slant Stack: A notion of Radon transform for data in a cartesian grid which is rapidly computible, algebraically exact, geometrically faithful and invertible*. Department of Statistics, Stanford University, 2001.

[15] Elias S Helou, Marcelo VW Zibetti, and Eduardo X Miqueles. Superiorization of incremental optimization algorithms for statistical tomographic image reconstruction. *Inverse Problems*, 33(4):044010, 2017.

[2] Download: `https://github.com/exmiqueles/raft.git`